

(12) UK Patent Application (19) GB (11) 2 362 062 (13) A

(43) Date of A Publication 07.11.2001

(21) Application No 0009114.0

(22) Date of Filing 13.04.2000

(71) Applicant(s)
3Com Corporation
(Incorporated in USA - Delaware)
5400 Bayfront Plaza, M/S 1308, Santa Clara,
California 95052-8145, United States of America

(72) Inventor(s)
David James Stevenson
Robert James Duncan
Alastair Hugh Chisholm
Ronan Francois Daniel Grandin
Neil William Gray

(74) Agent and/or Address for Service
Brookes Batchellor
102-108 Clerkenwell Road, LONDON, EC1M 5SA,
United Kingdom

(51) INT CL⁷
H04L 12/24 , H04Q 3/00

(52) UK CL (Edition S)
H4K KFMA

(56) Documents Cited
GB 2286317 A EP 0849910 A2 WO 97/31451 A1
US 6040834 A US 5615323 A US 5471399 A

(58) Field of Search
UK CL (Edition R) H4K KFM , H4P PEUL PFD
INT CL⁷ H04L 12/24 12/26 , H04Q 3/00
ONLINE: WPI, EPODOC, JAPIO.

(54) Abstract Title
Network management apparatus with graphical representation of monitored values

(57) A network management apparatus (3A) and method for processing network management data receives values for a monitored characteristic of a part of a network (1). Each time a new value for the monitored characteristic is received, a statistical value representing a level of the monitored characteristic over time is calculated. In one embodiment, the statistical value is a statistical average of received values over time, and in another embodiment the statistical value is the number of times the value of the monitored characteristic exceeds a predetermined threshold, for instance indicating high stress. The statistical value is then used to determine a graphical representation of the part of the network. For instance, an icon representing a device on the network has a size in proportion with the statistical value. Similarly, a line representing a link on the network has a thickness depending upon the statistical value. The determined graphical representation is then used in the display of a network map (17A)(Fig.3). Alternatively, the colour of an icon may change in dependence upon the statistical value.

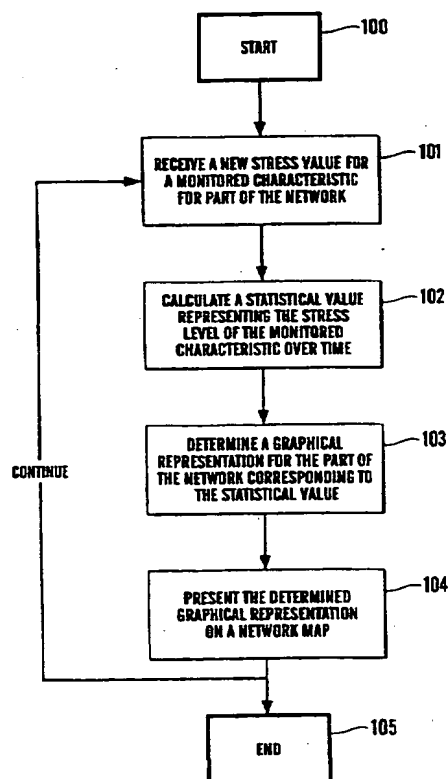


Fig.4

GB 2 362 062 A

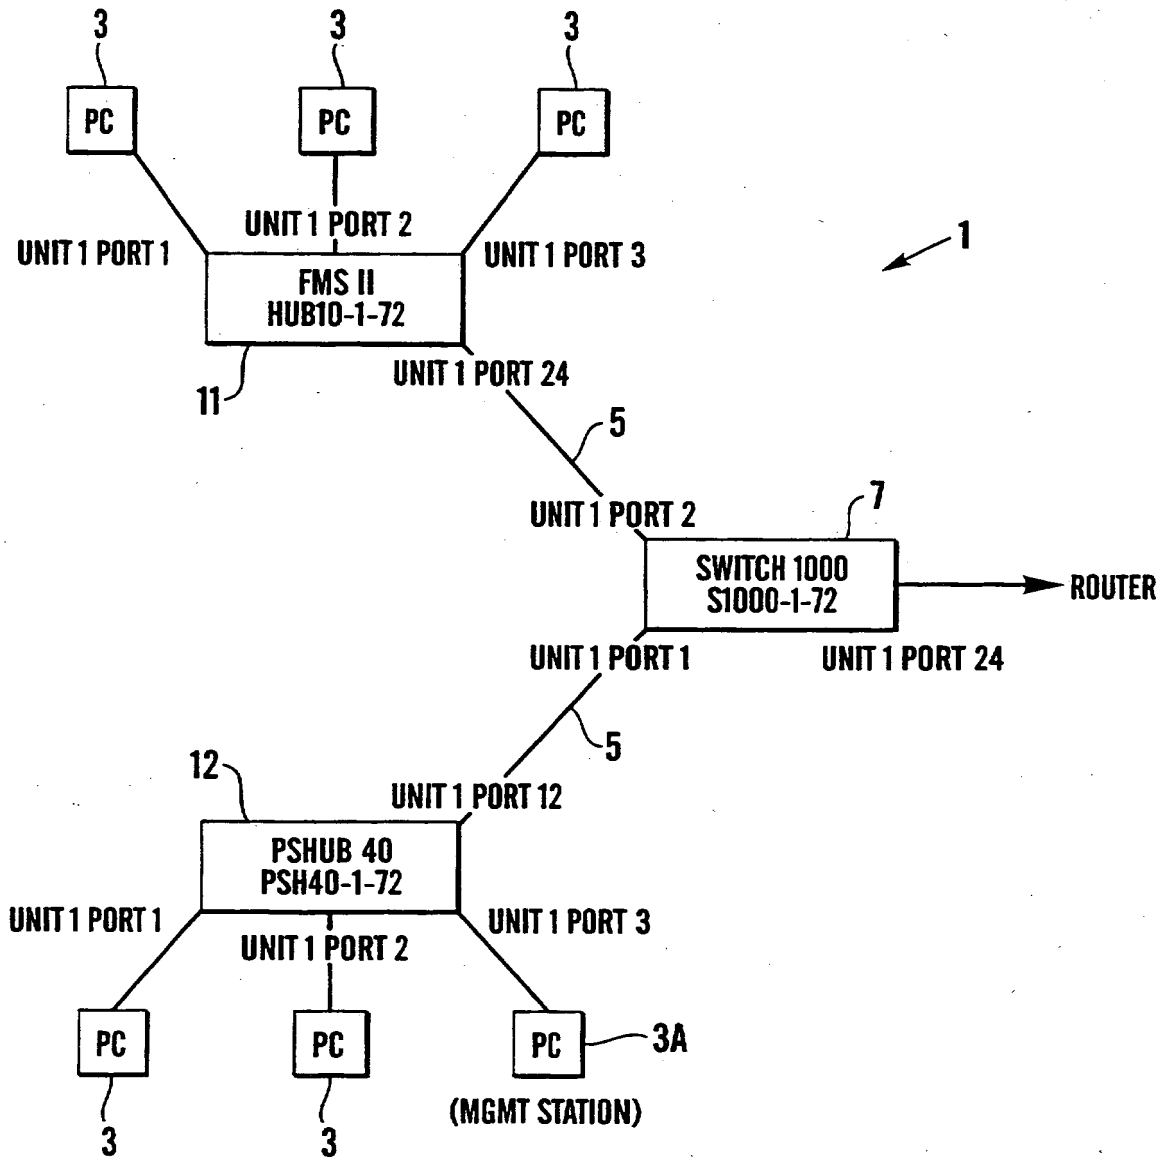



Fig. 1

Fig.2



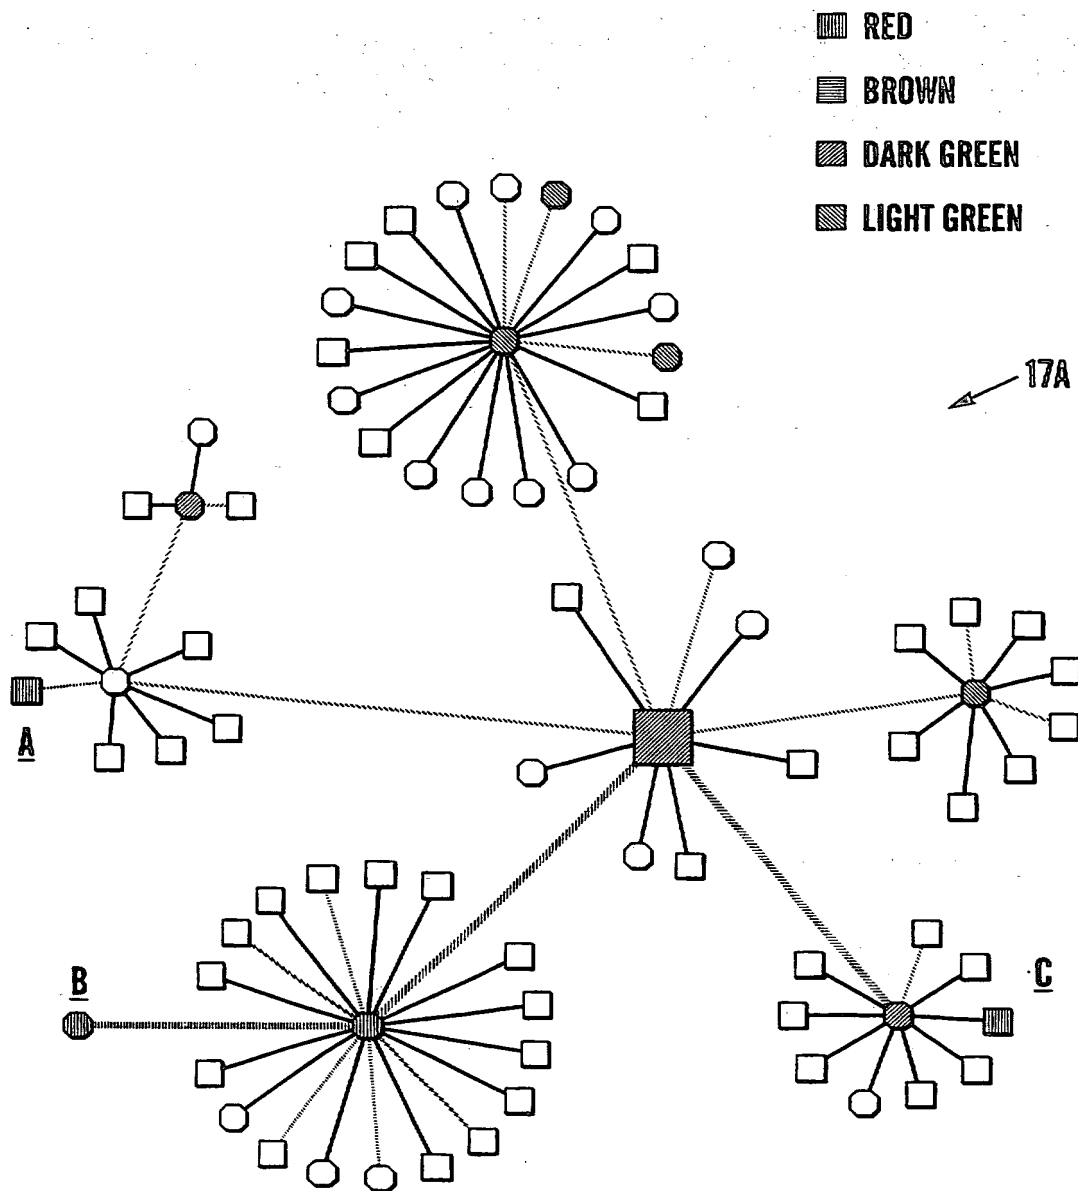
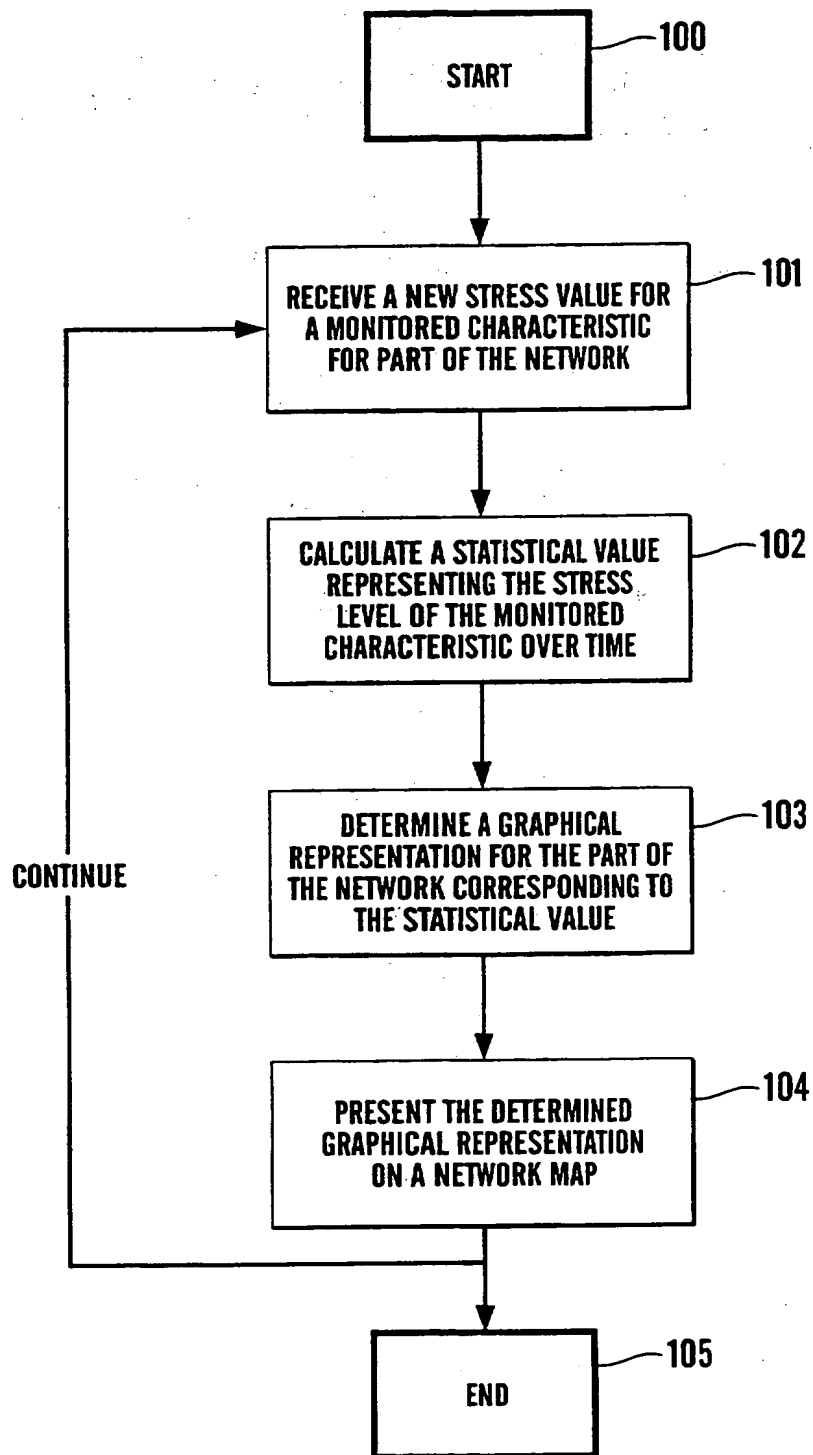


Fig. 3

*Fig.4*

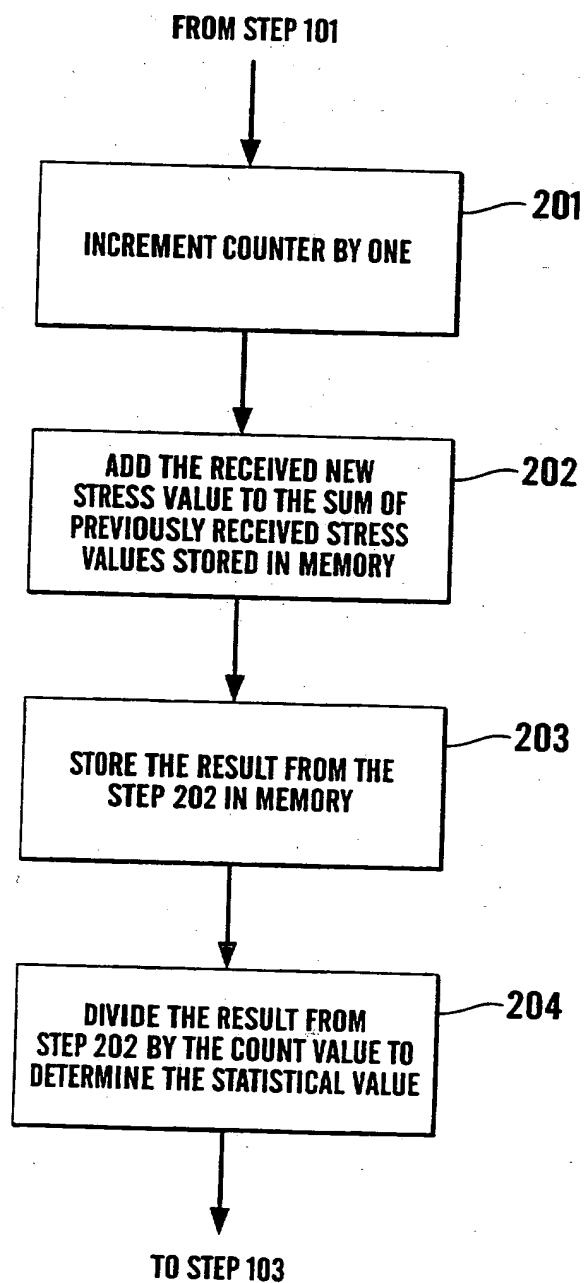


Fig.5

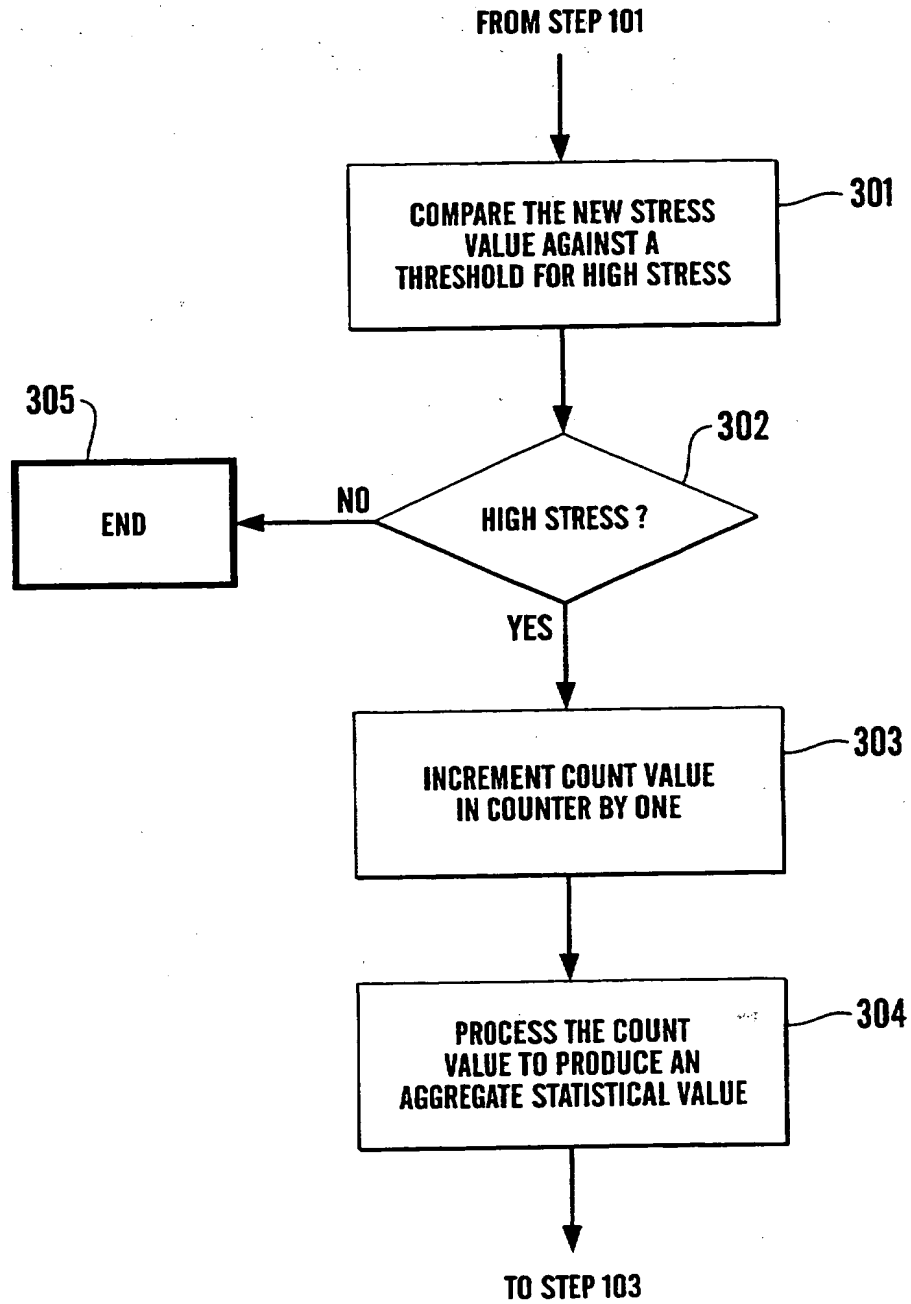


Fig.6

NETWORK MANAGEMENT APPARATUS AND METHOD FOR MONITORING STRESS OF A NETWORK

5 The present invention relates generally to an apparatus and method for the management of a network, and more particularly to a network management apparatus and method for use in monitoring the health or stress of a network.

10 The following description is concerned with a data communications network, and in particular a local area network (LAN) but has more widespread applicability to other managed communications systems including wide area networks (WANs) or wireless communications systems.

15 Networks typically comprise a plurality of computers, peripherals and other electronic devices capable of communicating with each other by sending and receiving data packets in accordance with a predefined network protocol. Each computer or other device on the network is connected by a port to the network media, which in the case of a LAN network may be coaxial cable, twisted pair cable or fibre optic cable. Each device on the network typically has hardware for media access control (MAC) with its own unique MAC address. Data packets are sent and received
20 in accordance with the MAC protocol (e.g. CSMA/CD protocol as defined by the standard IEEE 802.2, commonly known as Ethernet). Data packets transmitted using the MAC protocol identify the source MAC address (i.e. the MAC address of the device sending the data packet) and the destination MAC address (i.e. the MAC address of the device for which the data packet is destined) in the header of the data
25 packet.

30 A network is generally configured with core devices having a plurality of ports, which can be used to interconnect a plurality of media links on the network. Such devices include hubs, routers and switches which pass data packets received at one port to one or more of its other ports, depending upon the type of device. Such core devices can either be managed or unmanaged.

A managed device is capable of monitoring data packets passing through its ports. For example, a managed device can learn the physical or MAC addresses of the devices connected to its ports by monitoring the source address of data packets passing through the respective ports. Identified source addresses transmitted from a port of a managed network device, such as a router, hub or switch, are stored in a respective "address table" associated with the port

Managed devices additionally have the capability of communicating with each other using a management protocol such as the SNMP (Simple Network Management Protocol), as described in more detail below. Whilst the following description is concerned with the SNMP management protocol, the skilled person will appreciate that the invention is not limited to use with SNMP, but can be applied to managed networks using other network management protocols.

SNMP defines agents, managers and MIBs (where MIB is Management Information Base), as well as various predefined messages and commands for data communication. An agent is present in each managed network device and stores management data, responds to requests from the manager using the GETRESPONSE message and may send a TRAP message to the manager after sensing a predefined condition. A manager is present within the network management station of a network and automatically interrogates the agents of managed devices on the network using various SNMP commands such as GET and GETNEXT commands, to obtain information suitable for use by the network administrator, whose function is described below. A MIB is a managed "object" database which stores management data obtained by managed devices, accessible to agents for network management applications.

SNMP forms part of the TCP/IP protocol suite, which is a number of associated protocols developed for networks connected to the Internet also known as the World Wide Web.

It is becoming increasingly common for an individual, called the network administrator, to be responsible for network management, and his or her computer system or workstation is typically designated the network management station. The network management station incorporates the manager, as defined in the SNMP protocol, i.e. the necessary hardware, and software applications to retrieve data from MIBs by sending standard SNMP requests to the agents of managed devices on the network.

Network management software applications are known which can determine the topology of a network, i.e. the devices on the network and how they are linked together. In order to determine the network topology, the application retrieves data from the managed devices on the network, which data can provide information about the devices connected to the managed devices, for instance the aforementioned "address tables". MIB data can also be retrieved from managed devices to provide information about device type, device addresses and details about the links. Using such data, the application can usually determine the topology of the entire network.

An example of a known network management software application capable of determining network topology is the Transcend® Network Supervisor application available from 3Com Corporation of Santa Clara, California, USA.

A part of the network administrator's function is to identify and resolve problems occurring on the network, such as device or link malfunction or failure. In order to provide the network administrator with the necessary information to identify such problems, the network management application monitors the devices on the network. An example of such monitoring is described in co pending UK Patent Application No 9917993.9 entitled "Management System and Method for Monitoring Stress in a Network" in the name of the present applicant. In the system and method described in UK Patent Application No 9917993.9 the SNMP manager in the network management station requests the agents of managed network devices on the network device to retrieve selected MIB data indicative of device and link operation, and

performs tests for device activity and service availability. Such MIB data may relate to characteristics such as traffic activity or errors occurring at a particular port. Tests may include sending ICMP Ping requests to each device on the network, or sending selected requests for services such as SMTP, NFS and DNS to servers, and
5 monitoring the time taken to receive a response. The monitored parameters or characteristics are referred to herein as "stress metrics".

One way of presenting the monitored data to the network administrator is to provide an "Event list", which records occurrences of abnormally high values for
10 monitored stress metrics. In order to generate an Event list, the network management application compares, for each stress metric, the retrieved data or test result against a corresponding threshold level for the stress metric. The threshold level is the level above or below which performance is considered to be unacceptable. For
15 convenience, the following description assumes the threshold is the stress level which is deemed to be unacceptable so that stress values greater than or equal to the threshold level are unacceptable or "high". Default values for threshold levels of monitored stress metrics are typically preset by the application vendor, but may be adjusted by the network administrator.

Each time a threshold is exceeded, the application logs an "Event" in the
20 Event list or Event log. Each Event is entered on the list with relevant information such as the date and time of the Event, the identity of the device affected and the nature of the Event i.e. details of the stress metric value which exceeded the threshold. The network administrator can then review the Event list to identify problems on the
25 network.

A problem with this method of presenting management data to the network administrator is that the administrator has to piece together what has happened, based on the Events, to establish the underlying cause(s) of the Events. This may be
30 difficult for an inexperienced administrator and in any case is time consuming.

An alternative way of presenting the monitored data to the network administrator is to display a graphical representation of the network as a "network" map. Each device on the network is represented by an appropriate device icon and the links between devices are represented by lines connecting the device icons on the network map. The monitored stress values for each device and link are displayed in association with the corresponding icon on the network map (e.g. by clicking on a device icon to open a window showing the monitored values). The icons representing devices and lines representing links may be coloured to represent the stress levels. For example, the colour green may be used for low stress levels, and red may be used to represent high stress levels.

A problem with this method of presenting management data to the network administrator is that the network map represents the most recently monitored stress levels associated with the network. It does not reflect past performance and cannot provide clues to problems found previously, since the network map is continually updated as the devices are monitored. Thus, if a malfunctioning device or link happens to be working at the time the administrator views the network map, the administrator is not alerted to the problem.

It would be desirable to provide a network management system and method which addresses the aforementioned problems.

In accordance with a first aspect, the present invention provides a method for processing network management data comprising: receiving a new value for a monitored characteristic of a part of a network for which at least one previous value has been received; calculating a statistical value representing a level of the monitored characteristic over time, and determining a graphical representation of the part of the network corresponding to the statistical value.

The method can thus be used to produce a graphical representation for use in a network map to visibly distinguish parts of the network which have a history of high

stress levels from those which have a history of low stress levels.

5 In accordance with a second aspect, the present invention provides a computer readable medium carrying a computer program for carrying out the method of the first aspect of the present invention.

10 In accordance with a third aspect, the present invention provides a network management apparatus for operating the method of the first aspect of the present invention.

Embodiments of the present invention will now be described, by way of example, with reference to the accompanying drawings, in which:

15 Figure 1 is a block diagram of a typical network having a network management station in accordance with an embodiment of the present invention;

Figure 2 is a network map produced in accordance with a prior art network management software application;

20 Figure 3 is a network map generated in accordance with the present invention;

Figure 4 is a flow diagram showing the steps performed by a computer program implementing the method in accordance with a preferred embodiment of the present invention;

25 Figure 5 is a flow diagram showing substeps of one step in the program of Figure 4 according to a first embodiment of the present invention, and

30 Figure 6 is a flow diagram showing substeps of one step in the program of Figure 4 according to a second embodiment of the present invention.

Figure 1 shows a typical local area network 1. The network 1 operates various protocols within the TCP/IP protocol suite, as described above. The network 1 includes SNMP managed network devices including hubs 11 and 12, and switch 7 which are connected together by media links 5. End stations, in the form of personal computers (PCs) 3 are provided for users, including management station 3A for the network administrator. The PCs 3 are unmanaged network devices. Each PC 3 is connected to a core managed device, such as hub 11, which in turn is connected to switch 7. The hubs 11 and 12 and switch 7 are managed devices which communicate management information with management station 3A using the SNMP protocol. The agent within each core managed network device monitors data traffic passing through its ports and stores the data thereby obtained in an appropriate location in a MIB. Typically, the MIB data is represented in the form of "conceptual tables" (called "MIB tables") as is well known in the art. A typical managed device may implement a number of MIBs for network management.

An example of a MIB containing network management data is MIB-II (formerly MIB-I) defined by RFC 1213: Management Information Base for Network Management of TCP/IP Internets. In the network 1, hubs 11 and 12 and switch 7 store MIB-II data. Another MIB containing more complex management data is RMON (as defined in RFC 1757: Remote Network Monitoring Management Information Base) and in the network 1, only switch 7 stores RMON data.

Figure 2 shows a map 17 of the network 1 of Figure 1. Such a network map 17 is produced using a conventional network management software application, such as the Transcend® Network Supervisor application available from 3Com Corporation of Santa Clara, California, USA. Such a network management application "discovers" the network topology by interrogating the managed network devices i.e. hubs 11 and 12 and switch 7 on the network using the SNMP protocol. The data received is then processed to determine the number and type of network devices on the network, and how they are linked together. This information is then used to produce the network map which is displayed on a display screen 15 of the network

management station 3A when running the network management application or may be printed.

As shown in Figure 2, the network map 17 shows the network devices, each represented by an icon indicative of the type of network device, and links, each represented by a line joining two of the device icons. Associated with each icon is a label identifying each network device.

The network is typically discovered once by the SNMP manager in network management station 3A, in accordance with techniques such as those mentioned above, and the network map image is stored in memory associated with the management station 3A.

The network map 17 is used for monitoring the network 1. In particular, the network management station 3A periodically monitors the devices on the network by sending requests for selected MIB data from the managed network devices such as switch 7 and hubs 11 and 12 using the SNMP protocol. The retrieved management data is processed to determine the level of stress of the corresponding device, port of the device or part of the network as described in UK Patent Application No 9917993.9 mentioned above.

As in UK Patent Application No 9917993.9 a monitored characteristic indicative of stress is herein called a "stress metric". Examples of stress metrics include error rates for data frames and link utilisation, which are determined using MIB data, and Ping and service request response times which are determined by sending appropriate TCP/IP signals and monitoring the time taken to receive a response.

Typically, the network management station 3A monitors a plurality of stress metrics for a number of different parts of the network. The stress levels or values for the metrics are obtained by periodically requesting relevant MIB-II data from hubs 11

and 12 and switch 7, and RMON data from switch 7 and processing the received data, and by periodically polling all network devices using Ping or service requests and monitoring response times. For Ping and service requests, the management station 3A waits a predetermined period of time before it considers that a response has not been received and the request is "timed out".

The network management station 3A indicates the thus determined stress values in association with the network map 17. For instance, if no response is received to an IP Ping of an endstation 3, the network map will illustrate the endstation icon in red. Clicking on the icon will reveal further details i.e. that the endstation has a high level of stress because no response was received to the most recent IP Ping from the network management station 3A.

The network map 17 is continually updated, each time the network management station 3A receives a new stress value during its monitoring process. Thus, the network map 17 only represents the stress of the network 1 at a particular moment in time and does not represent past stress levels indicating past performance.

In accordance with a preferred embodiment of the present invention, the network management station 3A generates, either in addition to or in place of the conventional network map 17, a network map 17A representative of past and current stress levels. In particular, the network management station 3A determines an aggregate or average stress level for each monitored stress metric for each part of the network 1 and represents the average or aggregate stress levels on the network map.

Figure 3 shows a network map 17A of a more complex network representing the aggregate or average stress levels generated in accordance with the preferred embodiment. The network map 17A uses different colours for map icons and lines, different line thicknesses and different icon sizes to represent the aggregate or average stress level for the corresponding device or link.

An aggregate or average stress level may be calculated in a number of ways. In one example, each time a new stress value is determined for a monitored metric it is added to the sum of the previously recorded stress values for the same metric. The total is then divided by the number of recorded stress values to obtain a mean average for the stress value. A size and colour of the icon for the device or link being monitored is then selected to represent the average stress level. This method may be used to indicate areas of the network which are consistently and continually stressed.

Alternatively, in another example, the aggregate stress level may be determined based on the number of occasions that high stress is detected. For example, each time a high stress level is determined for a part of the network, i.e. the relevant threshold for the stress metric is exceeded, a corresponding counter is incremented. This count value is then used to represent the aggregate stress level on the network map, i.e. the size and colour of the icons for the device or link being monitored is selected to represent the counter value. This method may be used to indicate areas of the network which are frequently stressed, but which may not remain stressed for long periods of time (and hence which may have a low aggregate stress level when calculated using the method in the first example).

It will be appreciated that other methods of producing an average or aggregate stress level may be possible using statistical techniques, and combinations of the abovementioned and other methods may be employed.

The size and/or colour of icons for devices and the thickness and colour of the lines for links on the network map 17A may be chosen to visibly distinguish parts of the network which have high levels of stress from parts that have medium or low levels of stress. In the preferred embodiment, both colour and size are used to indicate stress levels of devices and links. In particular, small green icons are used for devices having a low aggregate stress, and thin green lines are used to represent links having low aggregate stress. Devices and links under medium stress are represented by a red/green mix and the corresponding icons are slightly larger and line thicknesses

slightly thicker than those representing low aggregate stress. Devices under high aggregate stress are represented by large red icons, and links under high stress are represented by thick red lines.

5 This may be achieved using an algorithm which maps an aggregated stress level onto a corresponding graphical representation of the device or link on the network map.

10 For example, consider an arbitrary numeric range for aggregated stress values for a link, such as 0 to 100, where 0 is minimum stress and 100 is maximum stress. An aggregate stress ("stress") for a link of 0 will map onto a line of minimum width ("min width" e.g. 1 pixel width) coloured "saturated green". An aggregate stress for a link of 100 will map onto a line of maximum width ("max width" e.g. 3 pixels) coloured "saturated red". Aggregate stress values between 0 and 100 will map
15 linearly according to the following equations:

$$\text{line width} = \text{min width} + ((\text{max width} - \text{min width}) / 100) * \text{stress}$$

$$\text{line colour} = \text{stress} / 100 * (\text{saturated red}) + (1 - (\text{stress} / 100)) * (\text{saturated green})$$

20 where the '+' operation indicates a mixing of the two colours.

A similar operation could be used to map the aggregate stress values for a device to a corresponding size and colour of icon.

25 The differences in size of the icons and lines and/or the use of colour in the network map 17A, enables the network administrator to readily identify areas of the network having high stress and to determine whether the parts of the network under high stress are related or isolated.

30 For instance, Figure 3 shows that there are a group of devices and links under high stress around node B. These problems are probably related. The high stress

levels recorded at nodes A and C are isolated. The network administrator immediately knows that he or she should concentrate on determining the root cause of the problems around node B.

5 Figure 4 shows the steps performed in a method according to a preferred embodiment of the present invention. It will be appreciated that the method is preferably implemented by a computer program operating in a network management station, but may equally be implemented in hardware. The following description relates to the method steps performed by a computer program.

10 The program starts at step 100, and at step 101 the program receives a new stress value for a monitored characteristic for a part of the network. The new stress value is typically received during the running of a conventional network management software application in which various stress metrics are periodically monitored, as
15 described above.

 In step 102, the program uses the new stress value to calculate a statistical value, typically an average or aggregate value representing the stress level of the monitored characteristic over time. Step 102 may be performed using the program
20 substeps illustrated in Figure 5, or alternatively using the substeps illustrated in Figure 6, described in detail below.

 In step 103, the program determines a graphical representation for the part of the network which corresponds to the statistical value determined in step 102. This
25 step may be performed using a look-up-table or algorithm as described above which maps the statistical value against a corresponding size, shape or colour of a graphical representation for various parts of the network.

 In step 104, the program presents, by display or printing, the determined
30 graphical representation of the part of the network on a network map. The program

ends at step 105 unless, in the meantime, it receives a new stress value, in which case the program continues from step 101.

Figure 5 shows the program substeps performed in step 102 of Figure 4 in accordance with a first example. In this first example, the program calculates an average stress value for the monitored characteristic. The stress value may be a normalised stress value, for instance determined as described in UK Patent Application No 9917993.9, or a stress value otherwise determined using raw data obtained during monitoring of the stress metric. The program stores the sum of previously received stress values for the monitored metric in memory and has a counter the value of which represents the number of stress values received for the monitored metric.

According to the substeps shown in Figure 5, when a new stress value is received by the program at step 101 of Figure 4, in step 201, the counter is incremented by one. The received stress value in step 102 is then added to the sum of previously received stress values for the monitored metric in step 202. This is achieved by retrieving the sum from memory, and adding it to the new stress value. The result is then stored in memory at step 203. The value stored in step 203 is then divided by the count value of the counter resulting from step 201, in step 204, thus determining the average stress value. The program then continues with step 103 of Figure 4.

Figure 6 shows the substeps performed by the program in step 102 of Figure 4, in accordance with a second example. In the second example, the value determined in step 102 is an aggregate value representing the number of times the stress value for the monitored characteristic of the relevant part of the network is high i.e. the stress value indicates unacceptable performance. As previously described, threshold values for each stress metric are predetermined, and stress values greater than or equal to the threshold are considered to represent high stress, and are considered to represent "Events".

Referring to Figure 6, after receiving a new stress value for the monitored characteristic in step 101 of Figure 4, in step 301 the program compares the new stress value against a predetermined threshold for high stress of the monitored characteristic.

5 In step 302, if the comparison in step 302 determines that the new stress value is less than the threshold, the program ends at step 305. Otherwise, if the comparison determines that the new stress value is greater than or equal to the high stress threshold, the program continues at step 303. In step 303, the program increments a counter by one to record that another Event has occurred. In step 304, the program
10 processes the count value to produce an aggregate value which is then used in step 103 of the program of Figure 4. The processing in step 304 may simply use the count value from the counter (which represents the number of times the threshold has been exceeded for the time the program has been running), or may divide the count value by a time period, so that the aggregate stress represents a rate of high stress Events,
15 rather than an absolute number of high stress Events.

As the skilled person will appreciate from the foregoing, in accordance with a preferred embodiment, the present invention is implemented in the form of a computer program within a network management software application, which is
20 provided on a computer readable medium. The computer readable medium may be a disk which can be loaded in a disk drive of the network management station 3A of the network of Figure 1. Alternatively, the computer readable medium may be a computer system carrying the website or other form of file server from which the computer program may be downloaded over the Internet or other network.

25 As the skilled person will appreciate, various modifications and changes may be made to the described embodiments. It is intended to include all such variations, modifications and equivalents which fall within the scope of the present invention as defined in the accompanying claims.

CLAIMS:

1. A method for processing network management data comprising
receiving a new value for a monitored characteristic of a part of a network for
5 which at least one previous value has been received;
calculating a statistical value representing a level of the monitored
characteristic over time, and
determining a graphical representation of the part of the network
corresponding to the statistical value.
10
2. A method as claimed in claim 1, wherein the monitored characteristic relates
to the stress of the part of the network.
3. A method as claimed in claim 1 or claim 2, wherein the statistical value is a
15 statistical average of received values for the monitored characteristic.
4. A method as claimed in claim 3, wherein the statistical average is determined
by the steps of:
adding together the new value and the at least one previously received values
20 for the monitored characteristic to produce a total; and
dividing the total by the number of received values for the monitored
characteristic.
5. A method as claimed in claim 4, further comprising storing the total in
25 memory.
6. A method as claimed in claim 4 or claim 5, further comprising, after the step
of receiving a new value and before the step of calculating, incrementing a counter by
one, wherein the count value of the counter represents the number of received values
30 for the monitored characteristic.

7. A method as claimed in claim 1 or claim 2, wherein the statistical value is calculated based on the number of times the received values for the monitored characteristic exceed a predetermined threshold value.

5 8. A method as claimed in claim 7, wherein the step of calculating comprises:
incrementing a counter if the new value exceeds the predetermined threshold,
wherein the count value of the counter represents the number of received values for
the monitored characteristic which exceed the predetermined threshold.

10 9. A method as claimed in claim 8, further comprising comparing the new value
with the predetermined threshold to determine if the new value exceeds the
predetermined threshold.

15 10. A method as claimed in claim 8 or claim 9, further comprising using the
counter value to determine the statistical value.

20 11. A method as claimed in any preceding claim, wherein the step of determining
a graphical representation comprises accessing a look-up-table which maps statistical
values for the monitored characteristic with corresponding graphical representations.

12. A method as claimed in any preceding claim, wherein the size of the graphical
representation changes in dependence on the statistical value.

25 13. A method as claimed in claim 12, wherein the size increases with an increase
in the statistical value.

14. A method as claimed in any preceding claim, wherein the colour of the
graphical representation changes in dependence on the statistical value.

30 15. A method as claimed in claim 14, wherein the colour is green for low
statistical values, and the colour is red for high statistical values.

16. A method for processing network management data substantially as hereinbefore described, with reference to, and as illustrated by Figures 4 and 5 or Figures 4 and 6 of the accompanying drawings.

5

17. A computer readable medium carrying a computer program for carrying out the method as claimed in any one of claims 1 to 16.

18. A network management apparatus for operating the method as claimed in any one of claims 1 to 16.

10



INVESTOR IN PEOPLE

Application No: GB 0009114.0
 Claims searched: 1-18

18

Examiner: Stephen Brown
 Date of search: 17 October 2000

Patents Act 1977 Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:
 UK Cl (Ed.R): H4P (PFD, PEUL), H4K (KFM).
 Int Cl (Ed.7): H04L: 12/24, 12/26, H04Q: 3/00.
 Other: Online: WPI, EPODOC. JAPIO.

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
Y	GB 2 286 317 A (Fujitsu) See especially fig.s 2, 4, 10, and page 4, line 36, to page 6, line 16, and page 8, lines 30-37.	1 at least
Y	EP 0 849 910 A2 (Northern Telecom) See especially the abstract, figure 18, and column 32, line 13.	1 at least
Y	WO 97/31451 A1 (MCI) See especially figure 5 and page 23, line 10, to page 28, line 7.	1 at least
Y	US 6 040 834 (Cisco) See especially figures 4-9, and column 6, line 17, to column 9, line 11, and claim 16.	1 at least
Y	US 5 615 323 (Concord) See especially column 3, line 55, to column 4, line 13, column 6, lines 10-59, and figure 2.	1 at least
Y	US 5 471 399 (Hitachi) See especially figures 13, 14, 16, 20 & 21, and column 8, lines 21-30, and column 9, line 58, to column 10, line 56.	1 at least

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.

THIS PAGE BLANK (USPTO)